Cycle Detection in Computation Tree Logic

Gaëlle Fontaine
Universidad de Chile
Santiago de Chile, Chile

Fabio Mogavero
University of Oxford
Oxford, UK

Aniello Murano
University of Naples
Naples, Italy

Giuseppe Perelli
University of Oxford
Oxford, UK

Loredana Sorrentino
University of Naples
Naples, Italy

Temporal logic is a very powerful formalism deeply investigated and used in formal system design and verification. Its application usually reduces to solving specific decision problems such as model checking and satisfiability. In these kind of problems, the solution often requires detecting some specific properties over cycles. For instance, this happens when using classic techniques based on automata, game-theory, SCC decomposition, and the like. Surprisingly, no temporal logics have been considered so far with the explicit ability of talking about cycles.

In this paper we introduce Cycle-CTL*, an extension of the classical branching-time temporal logic CTL* along with cycle quantifications in order to predicate over cycles. This logic turns out to be very expressive. Indeed, we prove that it strictly extends CTL* and is orthogonal to μCALCULUS. We also give an evidence of its usefulness by providing few examples involving non-regular properties.

We investigate the model checking problem for Cycle-CTL* and show that it is PSPACE-COMPLETE as for CTL*. We also study the satisfiability problem for the existential-cycle fragment of the logic and show that it is solvable in 2EXPTIME. This result makes use of an automata-theoretic approach along with novel *ad hoc* definitions of bisimulation and tree-like unwinding.

1 Introduction

Temporal logic is a suitable framework largely used in formal system verification [8, 10, 11, 25]. It allows to specify and reasoning in a rigorous manner about the temporal evolution of a system, without talking explicitly about the elapsing of time. Two fundamental decision problems involving temporal logics have been deeply investigated: *model checking* and *satisfiability*. The former, given a mathematical model of the system, such as a *Kripke* structure, asks whether it satisfies a temporal logic formula specifying its desired behavior. The latter, instead, checks whether the temporal logic specification is consistent and, thus, a corresponding system is feasible [10].

In several situations, reasoning about system correctness and, in particular, solving the above decision questions, reduces to detect precise cycle properties over the system model. For example, in the classical automata-theoretic approach there are settings in which the satisfiability question reduces to first build a Büchi automaton accepting all models of the formula and then to check for its non-emptiness [21]. The latter can be solved by looking for a "lasso", that is a path from the initial state to a final state belonging to a cycle [17, 21]. Similarly, if one uses a game-theory approach, solving the model checking or the satisfiability questions reduces to first construct a two-player game, such as a Büchi or a parity game [3, 13, 15, 21, 22, 29], and then check for the existence of a winning strategy for a designed player. The latter can be reduced to check whether it has the ability to confine the evolution of the game (a *play*) over some specific cycle over the arena, no matter how the other player behaves.

Depending on the view of the underlying nature of time, two types of temporal logics are mainly considered. In *linear-time temporal logics*, such as LTL [25], time is treated as if each moment in time has

D. Cantone and G. Delzanno (Eds.): Seventh Symposium on
Games, Automata, Logics and Formal Verification (GandALF'16)
EPTCS 226, 2016, pp. 164–177, doi:10.4204/EPTCS.226.12

© G. Fontaine et al.
This work is licensed under the
Creative Commons Attribution License.

a unique possible future. Conversely, in *branching-time temporal logics* such as CTL [8] and CTL* [12] each moment in time may split into various possible futures. Then, to express properties along one or all the possible futures we make use of existential and universal quantifiers. Noticeably, LTL is suitable to express path properties; CTL is more appropriate to express state-based property; finally, CTL* has the power to express combinations of path and state properties. In the years, these logics have been extended in a number of ways in order to express very complicated specification properties. Surprisingly, no temporal logic has been introduced so far to reason explicitly about cycles, despite their usefulness. In addition to the technical motivation mentioned above, there are often cases in which it is useful to distinguish between purely infinite behaviors, like those occurring in infinite-state systems, from regular infinite behaviors [6, 18]. Moreover, also in finite-state systems there are infinite behaviors that are not regular, like the *prompt*ones [19, 24], which we can distinguish by using our new concept of cycling path, as we show in an example later in the paper.

In this paper we introduce *Cycle*-CTL*, an extension of the classical logic CTL* along with the ability to predicate over cycles. For a *cycle* we mean a path that passes through its initial state infinitely often. Syntactically, Cycle-CTL* is obtained by enriching CTL* with two novel cycle quantifiers, namely the existential one $\mathtt{E}^{\circlearrowleft}$ and the universal one $\mathtt{A}^{\circlearrowleft}$. Note that Cycle-CTL* still uses the classical quantifiers E and A. Hence, we can use it to specify models whose behavior results as an opportune combination of standard paths and cycles. In particular, Cycle-CTL* can specify the existence of a lasso within a model.

We study the expressiveness of Cycle-CTL* and show that it is strictly more expressive than CTL* but orthogonal to μ-calculus. To give an evidence of the power and usefulness of the introduced logic, we provide some examples along the paper. Precisely, we first show how Cycle-CTL* can be used to reasoning, in a very natural way, about liveness properties restricted to cycles. Precisely, we show how to specify that some designed property recurrently occurs in the starting state of a cycle. As another example, we show the ability of the logic to handle non-regular properties such as the "*prompt-parity condition*" [24]. In temporal logic, we can specify properties that will eventually hold, but this gives no bound on the moment they will occur. Prompt temporal logics and games have been deeply investigated in order to restrict reasoning about properties that only occur in bounded time [1, 4, 7, 20, 24].

We investigate both the model checking and the satisfiability questions for Cycle-CTL* and provide some automata-based solutions. For the model checking question we provide a PSPACE upper-bound by opportunely extending the classical approach that is used for CTL* [21]. Specifically, we add a machinery consisting of an appropriate Büchi automaton that checks in parallel whether a path is a cycle and satisfies a required formula. Concerning the satisfiability question, we introduce instead a novel approach that makes use of two-way automata [26]. These automata, largely investigate and used in formal verification [5, 14, 18], allow to traverse trees both in forward and backward. The reason why we cannot use and extend the classical approach provided for CTL* (see [21]) resides on the fact that such an approach makes strongly use of some positive properties that hold for CTL*, among the others the tree- and the finite-model ones. Unluckily and unsurprisingly, due to the ability in Cycle-CTL* to force (and even more to forbid) the existence of cycles, we lose in this logic both these properties. This requires the introduction of novel and *ad hoc* definitions of bisimulation and tree-like unwinding to be used along with the automata-based approach. In particular, two-way tree automata are used to collect all tree representations of such tree-like unwinding structures. By means of this machinery we show that the satisfiability question for the full logic is 3EXPTIME. We also investigate the satisfiability of the existential-cycle fragment and show that it is solvable in 2EXPTIME, thus it is not harder than CTL* in complexity. Such fragment is simply obtained by forbidding the use of the universal cycle quantifier $\mathtt{A}^{\circlearrowleft}$ and by only allowing negations over atomic propositions. Note that this fragment still admits the classical E and A quantifiers, as well as it strictly subsumes CTL*.

2 Computation-Tree Logic with Cycle Detection

In this section we introduce and discuss the syntax and semantics of Cycle-CTL* (CTL$^\star_\circlearrowleft$, for short) and discuss some interesting problems that can be expressed in our logic.

Models We first provide the definition of the underlying model for our Cycle-CTL*.

Definition 1 (Kripke Structure)**.** *A* Kripke structure *(*KS*, for short) [16] over a finite set of* atomic propositions AP *is a tuple* $\mathscr{K} \triangleq \langle \mathrm{AP}, \mathrm{W}, R, \mathsf{L}, w_I \rangle$*, where* W *is an enumerable non-empty set of* worlds*,* $w_0 \in \mathrm{W}$ *is a designated* initial world*,* $R \subseteq \mathrm{W} \times \mathrm{W}$ *is a left-total* transition relation*, and* $\mathsf{L} : \mathrm{W} \mapsto 2^{\mathrm{AP}}$ *is a* labeling function *mapping each world to the set of atomic propositions true in that world.*

A *path* in $\mathscr{K}$ is an infinite sequence of worlds $\pi \in \mathrm{Pth} \subseteq \mathrm{W}^\omega$ such that, for all $i \in \mathbb{N}$, it holds that $((\pi)_i, (\pi)_{i+1}) \in R$. We denote by $\mathsf{fst}(\pi) \triangleq \pi_0$ and $(\pi)_i \triangleq \pi_i$ the first and i-th element of π. For a path π, we say that π is a *cycle* if, for all $i \in \mathbb{N}$, there exists $j \in \mathbb{N}$, with $j > i$, such that $(\pi)_j = \mathsf{fst}(\pi)$. For a given path π, we denote by $\mathsf{L}(\pi)$ the sequence γ in $(2^{\mathrm{AP}})^\omega$ such that $(\gamma)_i = \mathsf{L}(\pi_i)$ for all $i \in \mathbb{N}$. Moreover, $(\pi)_{\leq i} \triangleq \pi_0 \cdots \pi_i$ and $(\pi)_{\geq i} \triangleq \pi_i \cdot \pi_{i+1} \cdots$ represent the prefix up to and the suffix from position i of π. Prefixes of a path are also called *tracks* and denoted by $\rho \in \mathrm{Trk} \subseteq \mathrm{W}^+$. We also denote by $\mathsf{lst}(\rho)$ the last element occurring in the track ρ. Finally, all the definitions given above for paths naturally apply to tracks.

By $\mathrm{Trk}(w)$ and $\mathrm{Pth}(w)$ we denote the set of tracks and paths starting from w, respectively. By Cyc and $\mathrm{Cyc}(w)$ we denote the set of cycles and the set of cycles starting from w, respectively. Intuitively, tracks and paths of a KS $\mathscr{K}$ are legal sequences, either finite or infinite, of reachable worlds that can be seen as partial or complete descriptions of possible *computations* of the system modelled by $\mathscr{K}$.

For a pair $(w_1, w_2) \in R$, we say that w_2 is an R-successor of w_1. Note that in case R is a function, then each world w has only one R-successor. This implies that, starting from the initial world w_0, there is a unique legal path. Such structures are called LTL *models*.

Syntax CTL$^\star_\circlearrowleft$ extends CTL* [9] by means of two additional path operators, $\mathtt{E}^\circlearrowleft\psi$ and $\mathtt{A}^\circlearrowleft\psi$, which respectively read as "there exists a cycle path satisfying ψ" and "for all cycle paths ψ holds". As for CTL*, the syntax includes path-formulas, expressing properties over sequences of words, and state-formulas, expressing properties over a single word. State and path formulas are defined by mutual induction as follows.

Definition 2 (CTL$^\star_\circlearrowleft$ syntax)**.** CTL$^\star_\circlearrowleft$ *formulas are inductively built from a set of atomic propositions* AP*, by using the following grammar, where* $p \in \mathrm{AP}$*:*

$$\phi := p \mid \neg\phi \mid \phi \wedge \phi \mid \phi \vee \phi \mid \mathtt{E}\psi \mid \mathtt{A}\psi \mid \mathtt{E}^\circlearrowleft\psi \mid \mathtt{A}^\circlearrowleft\psi$$
$$\psi := \phi \mid \neg\psi \mid \psi \wedge \psi \mid \psi \vee \psi \mid \mathtt{X}\psi \mid \psi\mathtt{U}\psi$$

All the formulas generated by a ϕ-rule are called *state-formulas*, while the formulas generated by a ψ-rule are called *path-formulas*. By $\mathsf{sub}(\varphi)$ we denote the set of all subformulas of φ, and by $\mathsf{sub}_s(\varphi)$ we denote the set of state subformulas of φ.

Semantics The semantics for CTL$^\star_\circlearrowleft$ is defined *w.r.t.* Kripke structures. It extends the one for CTL*, with the addition of two new definitions for two cycle path quantifiers.

Definition 3. *The semantics of* CTL$^\star_\circlearrowleft$ *formulas is recursively defined as follows. For a Kripke structure* $\mathscr{K}$*, a world* w*, a path* π *and a natural number* $i \in \mathbb{N}$*, we have that:*

- *For all state formulas* ϕ*,* ϕ_1*, and* ϕ_2*:*

- $\mathscr{K},w \models p$ *if* $p \in \mathsf{L}(w)$;
- $\mathscr{K},w \models \neg\phi$ *if* $\mathscr{K},w \not\models \phi$;
- $\mathscr{K},w \models \phi_1 \wedge \phi_2$ *if both* $\mathscr{K},w \models \phi_1$ *and* $\mathscr{K},w \models \phi_2$;
- $\mathscr{K},w \models \phi_1 \vee \phi_2$ *if either* $\mathscr{K},w \models \phi_1$ *or* $\mathscr{K},w \models \phi_2$;
- $\mathscr{K},w \models \mathtt{E}\psi$ *if there exists a path* π *in* $\mathrm{Pth}(w)$ *such that* $\mathscr{K},\pi,0 \models \psi$;
- $\mathscr{K},w \models \mathtt{A}\psi$ *if, for all paths* π *in* $\mathrm{Pth}(w)$, *it holds that* $\mathscr{K},\pi,0 \models \psi$
- $\mathscr{K},w \models \mathtt{E}^{\circlearrowleft}\psi$ *if there exists a path* π *in* $\mathrm{Cyc}(w)$ *and* $\mathscr{K},\pi,0 \models \psi$;
- $\mathscr{K},w \models \mathtt{A}^{\circlearrowleft}\psi$ *if, for all paths* π *in* $\mathrm{Cyc}(w)$, *it holds that* $\mathscr{K},\pi,0 \models \psi$.

• *For path formulas* ϕ, ψ, ψ_1, *and* ψ_2:

- $\mathscr{K},\pi,i \models \phi$ *if* $\mathscr{K},(\pi)_i \models \phi$;
- $\mathscr{K},\pi,i \models \neg\psi$ *if* $\mathscr{K},\pi,i \not\models \psi$;
- $\mathscr{K},\pi,i \models \psi_1 \wedge \psi_2$ *if both* $\mathscr{K},\pi,i \models \psi_1$ *and* $\mathscr{K},w \models \psi_2$;
- $\mathscr{K},\pi,i \models \psi_1 \vee \psi_2$ *if either* $\mathscr{K},\pi,i \models \psi_1$ *or* $\mathscr{K},w \models \psi_2$;
- $\mathscr{K},\pi,i \models \mathtt{X}\psi$ *if* $\mathscr{K},\pi,i+1 \models \psi$;
- $\mathscr{K},\pi,i \models \psi_1 \mathtt{U} \psi_2$ *if there exists* $k \in \mathbb{N}$ *such that* $\mathscr{K},\pi,i+k \models \psi_2$ *and* $\mathscr{K},\pi,i+j \models \psi_1$, *for all* $j \in [0,k[$;

We say that π *satisfies the path formula* ϕ *over* $\mathscr{K}$, *and write* $\mathscr{K},\pi \models \phi$, *if* $\mathscr{K},\pi,0 \models \phi$. *Also, we say that* $\mathscr{K}$ *satisfies the state formula* φ, *and write* $\mathscr{K} \models \varphi$, *if* $\mathscr{K},w_I \models \varphi$.

Examples In this section, we provide some properties that are expressible with $\mathrm{CTL}^{\star}_{\circlearrowleft}$.

Assume that there is a system composed by two processes, requesting to access a resource, and a scheduler, releasing such resource in a fair way, *i.e.*, the resource is never used by the two processes at the same time. Every time the scheduler grants the resource to process i, such resource is exclusively used by process i until the system goes back to the decision point, that is, the state in which the scheduler released the resource. We denote by $\mathtt{dec}$ the atomic proposition labeling the states that are decision points (that is, the moment where the scheduler makes a decision) and by $\mathtt{res}_1$, $\mathtt{res}_2$ the atomic propositions representing the fact that the resource is released to processes 1 and 2, respectively. The above described situation can be expressed with the $\mathrm{CTL}^{\star}_{\circlearrowleft}$ formula $\varphi_i = \mathtt{E}^{\circlearrowleft}((\mathtt{dec} \wedge \neg\mathtt{res}_i \wedge \mathtt{G}\neg\mathtt{res}_{1-i}) \rightarrow \mathtt{F}\,\mathtt{res}_i)$, for $i \in \{1,2\}$. Note that in that formula, the use of the cycle operator is crucial as it allows us to loop at the decision point. As another example, we can also force the system to satisfy the mutual exclusion property in each possible decision point by means of the formula $\mathtt{AG}(\mathtt{dec} \rightarrow \varphi_1 \wedge \varphi_2)$. Finally, note that, since the system is required to loop on a decision point from which it is possible to release the resource for either process 1 or process 2, this automatically implies the existence of an infinite path which is able to satisfy the fairness condition, which is expressible in $\mathrm{CTL}^{\star}$ by means of the formula $\psi = \mathtt{E}(\mathtt{GFres}_1 \wedge \mathtt{GFres}_2)$. In other words, we have that $\varphi_1 \wedge \varphi_2 \rightarrow \psi$ is a valid $\mathrm{CTL}^{\star}_{\circlearrowleft}$ formula.

We now discuss another example involving prompt parity games, introduced in [24].

A Parity Game is a tuple of the form $\mathscr{P} = \langle \mathrm{V}, \mathrm{V}_0, \mathrm{V}_1, E, \mathsf{p}, v_0 \rangle$ where V is a nonempty finite set of states of the game, partitioned into V_0 and V_1, being the set belonging to Player 0 and Player 1, respectively, $E \subseteq \mathrm{V} \times \mathrm{V}$ is an edge relation, $\mathsf{p} : \mathrm{V} \rightarrow \mathbb{N}$ is a priority labeling function, assigning a natural number to each state, and $v_0 \in \mathrm{V}$ is a designated initial state. The game is played starting from v_0. At each state v of the game, if $v \in \mathrm{V}_i$, then Player i move to an E-successor of v. Such operation induces an infinite path π over V called play and then, by means of the function p, we also consider the infinite path $\mathsf{p}(\pi)$. Every occurrence of an odd priority on $\mathsf{p}(\pi)$ is called request. For any request, the successive

occurrence of an even and greater priority is its response. We say that Player 0 wins the play π under the parity condition if every request occurring infinitely often is responded. Moreover, we say that Player 0 wins the play π under the prompt parity condition if there exists a natural number n such that each request occurring infinitely often is responded in less than n steps. For both the cases above, we say that Player 1 wins the game iff Player 0 does not win. A strategy for Player i is a function $\mathsf{f}_i : \mathrm{V}^* \cdot \mathrm{V}_i \to \mathrm{V}$ assigning an E-successor to each partial (finite) path of the game. Clearly, a pair of strategies f_0 and f_1 determines a unique path and therefore, the winner. A strategy f_i is *positional* if, for all partial paths ρ and ρ', with $\mathsf{lst}(\rho) = \mathsf{lst}(\rho') \in \mathrm{V}_i$, it holds that $\mathsf{f}_i(\rho) = \mathsf{f}_i(\rho')$.

Let $\mathscr{P}$ be a parity game and f_0 be a positional strategy for Player 0. By projecting the strategy on the arena, we obtain a KS $\mathscr{K}_{\mathscr{P},\mathsf{f}_0} = \langle \mathrm{AP}, \mathrm{W}, R, \mathsf{L}, w_I \rangle$ defined as follows: $\mathrm{AP} = \mathsf{rng}(\mathsf{p}) = [0,n]$ [1], for some $n \in \mathbb{N}$, $\mathrm{W} = \mathrm{V}$, $R = \mathsf{f}_0 \cup E \cap (\mathrm{V}_1 \times \mathrm{V})$, $\mathsf{L}(w) = \{\mathsf{p}(w)\}$, for all $w \in \mathrm{W}$, and $w_I = v_0$. We can express that f_0 is winning for Player 0 by means of the formula $\varphi^{par} = \mathtt{A}(\bigvee_{k \equiv_2 0}(\mathtt{GF}k \wedge \bigwedge_{l \geq k, l \equiv_2 1} \mathtt{FG}\neg l))$. Indeed, the formula says that, for all possible paths, there exists an even priority k occurring infinitely often such that each odd priority l greater than k occurs finitely many times. Hence, we have that f_0 is winning over $\mathscr{P}$ iff $\mathscr{K}_{\mathscr{P},\mathsf{f}_0}, v_0 \models \varphi^{par}$.

In addition to this, we can express the existence of a path violating the prompt condition by means of the formula $\varphi^{npmt} = \bigvee_{n \equiv_2 0} \mathtt{E}(\bigvee_{k<n, k \equiv_2 1}(\mathtt{GF}k \wedge \mathtt{G}(k \to (\bigwedge_{l \geq k, l \equiv_2 0} \mathtt{G}\neg l)\mathtt{U}(\mathtt{E}^{\circlearrowleft} \bigwedge_{l \geq k, l \equiv_2 0} \mathtt{G}\neg l)))))$. At this point, the formula $\varphi^{par} \to \varphi^{npmt}$ is able to express the existence of a winning strategy for Player 1 under the prompt parity condition.

3 Model-Theoretic Properties

This section consists of two parts. First, we present invariance properties of $\mathrm{CTL}^{\star}_{\circlearrowleft}$. As trees do not contain any cycle and bisimulation do not preserve cycles, it does not come at a surprise that $\mathrm{CTL}^{\star}_{\circlearrowleft}$ is not invariant under bisimulation and does not have a tree-model property. Therefore, we introduce a new notion of bisimulation namely *cycle-bisimulation*, which takes cycles into account. We prove that $\mathrm{CTL}^{\star}_{\circlearrowleft}$ is invariant under cycle-bisimulation. Using that property, we show that $\mathrm{CTL}^{\star}_{\circlearrowleft}$ has a tree-like model property.

In the second part of the section, we investigate the expressive power of $\mathrm{CTL}^{\star}_{\circlearrowleft}$. We show that $\mathrm{CTL}^{\star}_{\circlearrowleft}$ strictly extends $\mathrm{CTL}^{\star}$ and is orthogonal to the μCALCULUS.

Invariance Properties We start by establishing that $\mathrm{CTL}^{\star}_{\circlearrowleft}$ is not invariant under bisimulation and does not have a tree-model or finite-model property.

Theorem 1 ($\mathrm{CTL}^{\star}_{\circlearrowleft}$ Negative Model Properties)**.** $\mathrm{CTL}^{\star}_{\circlearrowleft}$ *has neither the finite-model property, nor the tree-model property. It is also not invariant under bisimulation.*

Proof. Consider the formula $\varphi_1 = \mathtt{AG}\neg\mathtt{E}^{\circlearrowleft}\top$ stating that all paths starting from the initial state, do not contain any cycle. This formula is satisfiable. However, since the transition relation is such that each state has a successor, φ_1 can only be true in an infinite model.

Consider now the formula $\varphi_2 = \mathtt{E}^{\circlearrowleft}\top$. It is true in a model iff its initial state is the first point of a cycle. So φ_2 is satisfiable but is never true at the root of a tree. Hence, $\mathrm{CTL}^{\star}_{\circlearrowleft}$ does not have the tree-model property and thus, is not invariant under bisimulation. □

[1] *W.l.o.g.*, we can assume that the range of a priority function is an initial segment of $\mathbb{N}$.

Definition 4 (Bisimulation). *Let $\mathscr{K}_1 = \langle \mathrm{AP}_1, \mathrm{W}_1, R_1, \mathsf{L}_1, w_I^1 \rangle$ and $\mathscr{K}_2 = \langle \mathrm{AP}_2, \mathrm{W}_2, R_2, \mathsf{L}_2, w_I^2 \rangle$ be two Kripke structures. Then, a relation $B \subseteq \mathrm{W}_1 \times \mathrm{W}_2$ is a* cycle-bisimulation relation *if the following hold:*

1. *(w_I^1, w_I^2) belongs to B;*
2. *for all $w_1 \in \mathrm{W}_1$ and $w_2 \in \mathrm{W}_2$, if (w_1, w_2) belongs to B, then:*
   (a) *$\mathsf{L}_1(w_1) = \mathsf{L}_2(w_2)$;*
   (b) *for all $v_1 \in \mathrm{W}_1$ such that $(w_1, v_1) \in R_1$, there is $v_2 \in \mathrm{W}_2$ such that $(w_2, v_2) \in R_2$ and $(v_1, v_2) \in B$;*
   (c) *for all $v_2 \in \mathrm{W}_2$ such that $(w_2, v_2) \in R_2$, there is $v_1 \in \mathrm{W}_1$ such that $(w_1, v_1) \in R_1$ and $(v_1, v_2) \in B$;*
   (d) *for all cycles π_1 with beginning state w_1, there is a cycle π_2 with beginning state w_2 such that for all $i \in \mathbb{N}$, the pair $((\pi_1)_i, (\pi_2)_i)$ belongs to B,*
   (e) *for all cycles π_2 with beginning state w_2 there is a cycle π_1 with beginning state w_1 such that for all $i \in \mathbb{N}$, the pair $((\pi_1)_i, (\pi_2)_i)$ belongs to B.*

We say that $\mathscr{K}_1$ and $\mathscr{K}_2$ are cycle-bisimilar w.r.t. *a relation $B \subseteq \mathrm{W}_1 \times \mathrm{W}_2$ if B is a cycle bisimulation. Moreover, two paths π_1 and π_2 are* bisimilar w.r.t. *a cycle-bisimulation B if for all $i \in \mathbb{N}$, the pair $((\pi_1)_i, (\pi_2)_i)$ belongs to B.*

The notion of cycle-bisimulation is quite intuitive. While the usual definition of a bisimulation allows us to "mimic" the transition relation from one model to the other, a cycle-bisimulation also ensures that we can "mimic" cycles from one model to the other.

As a remark, the cycle-bisimulation notion is interesting by itself, as it gives rise to a new notion of equivalence among structures, that might lead to model-reduction characterization of the logic. We plan to investigate this aspect in a future work.

Theorem 2 (Invariance under bisimulation). *$\mathrm{CTL}^\star_\circlearrowleft$ is invariant under cycle-bisimulation.*

Using the invariance under cycle-bisimulation, we establish a tree-like model property for $\mathrm{CTL}^\star_\circlearrowleft$. Intuitively, the tree-model property for $\mathrm{CTL}^\star_\circlearrowleft$ fails as trees do not admit any cycle. Hence, the idea is to consider structures obtained by adding some restricted form of cycles over trees. We call those structures *trees with back edges* and they are defined as follows.

Definition 5. *A Kripke model $\mathscr{K} = \langle \mathrm{AP}, \mathrm{W}, R, \mathsf{L}, w_I \rangle$ is a* tree with back edges *if there are a Kripke model $\mathscr{T}_0 = (\mathrm{AP}, \mathrm{W}, R_0, \mathsf{L}, w_I)$ and a partial map $\mathsf{f} : \mathrm{W} \to \mathrm{W}$ such that*

(i) (W, R_0) is a tree with root w_I over the alphabet [2] *AP,*

(ii) R is equal to $R_0 \cup \{(w, \mathsf{f}(w)) : w$ belongs to the domain of $\mathsf{f}\}$,

(iii) for all $w \in \mathrm{W}$, $\mathsf{f}(w)$ is an ancestor of w,

(iv) for all $w_1, w_2 \in \mathrm{W}$, if $\mathsf{f}(w_1)$ is defined, $(\mathsf{f}(w_2), w_1), (w_1, w_2) \in {R_0}^+$ [3]*, then $\mathsf{f}(w_1) = \mathsf{f}(w_2)$.*

We say that $(\mathscr{T}_0, \mathsf{f})$ is a tree decomposition *of $\mathscr{K}$, where $\mathscr{T}_0$ is the* associated tree *and f is the* back-edge map. *If a pair (w, v) belongs to R_0, we say that (w, v) is associated with a* forward edge*, while if $v = f(w)$, the pair (w, v) is associated with a* back edge*.*

Note that if for every pair (w, v) in R we know whether (w, v) is associated with a forward or back edge, then this uniquely defines a tree decomposition.

[2] The relation R_0 is the child relation of the tree.

Intuitively, a tree with back edges is a structure obtained from a tree by adding edges (called back edges) from some nodes to their ancestors. More precisely, we add a back edge from each node w in the domain of f to its image $\mathsf{f}(w)$. Such back edges need to satisfy two conditions. First, each node must admit at most one outgoing back edge. The second condition (condition (iv)) is a bit less intuitive. It requires that the partial map f preserves the ancestor relation, and, in addition, that the back edges cannot "superpose", that is, in a tree back edges never cross each other.

We prove now the tree-like model property and show that each satisfiable formula of $\mathrm{CTL}^{\star}_{\circlearrowleft}$ is satisfiable in a tree with back edges. More specifically, given a Kripke model $\mathscr{K}$, we show how to define a tree with back edges $\mathscr{U}_{\mathscr{K}}$ such that $\mathscr{K}$ and $\mathscr{U}_{\mathscr{K}}$ are cycle-bisimilar. Together with Theorem 2, this implies that each satisfiable formula of $\mathrm{CTL}^{\star}_{\circlearrowleft}$, is satisfiable in a tree with back edges. Before defining $\mathscr{U}_{\mathscr{K}}$, we need to introduce two preliminaries notions: the projection map and the initial cycle state.

Let $\mathscr{K} = \langle \mathrm{AP}, \mathrm{W}, R, \mathsf{L}, w_I \rangle$ be a Kripke model and consider two constants $\mathtt{new}$ and $\mathtt{cycle}$. We define the *projection map* $\mathtt{pr} : (\mathrm{W} \times \{\mathtt{new}, \mathtt{cycle}\})^* \to \mathrm{W}$ as the unique surjective map such that:

$\mathtt{pr}(\varepsilon) = w_I$ and for all $w^\bullet \neq \varepsilon$, we have $\mathtt{pr}(w^\bullet) = w$, where $\mathsf{lst}(w^\bullet) = (w, \alpha)$ and $\alpha \in \{\mathtt{new}, \mathtt{cycle}\}$.

Given a state $w^\bullet \in (\mathrm{W} \times \{\mathtt{new}, \mathtt{cycle}\})^*$, we say that $w^\bullet$ admits a sequence $v^\bullet$ as an *initial cycle state* if there is a sequence $v_1 \dots v_k$ (where $k \geq 2$) such that $w^\bullet$ is equal to $v^\bullet\,(v_1, \mathtt{new})(v_2, \mathtt{cycle}) \dots (v_k, \mathtt{cycle})$. Given a sequence $w^\bullet \in (\mathrm{W} \times \{\mathtt{new}, \mathtt{cycle}\})^*$ such that $\mathsf{lst}(w^\bullet) = (w, \alpha)$, we say that $w^\bullet$ is labeled by w and α. Intuitively, the initial cycle state of a given state $w^\bullet$ is simply the parent of the closest ancestor of $w^\bullet$ that is labeled by $\mathtt{new}$. Note that a state admits at most one initial cycle state. We are now ready to define $\mathscr{U}_{\mathscr{K}}$.

Definition 6. *Given a Kripke model $\mathscr{K} = \langle \mathrm{AP}, \mathrm{W}, R, \mathsf{L}, w_I \rangle$, we define the* tree-like unwinding *$\mathscr{U}_{\mathscr{K}} = \langle \mathrm{AP}^\bullet, \mathrm{W}^\bullet, R^\bullet, \mathsf{L}^\bullet, w\bullet_I \rangle$ of $\mathscr{K}$ in the following way:*

- $\mathrm{W}^\bullet = (\mathrm{W} \times \{\mathtt{new}, \mathtt{cycle}\})^*$;
- $w_I^\bullet = \varepsilon$;
- *for all $w^\bullet \in \mathrm{W}^\bullet$, we have $\mathsf{L}^\bullet(w^\bullet) = \mathsf{L}(\mathtt{pr}(w^\bullet))$;*
- *for all $w^\bullet \in \mathrm{W}^\bullet$ and for all $(\mathtt{pr}(w^\bullet), v) \in R$:*
  - *the pair $(w^\bullet, w^\bullet(v, \mathtt{new}))$ belongs to $R^\bullet$ and is associated with a forward edge;*
  - *if $w^\bullet$ admits an initial cycle state $u^\bullet$ such that $\mathtt{pr}(u^\bullet) \neq v$, then the pair $(w^\bullet, w^\bullet(v, \mathtt{cycle}))$ belongs to $R^\bullet$ and is associated with a forward edge;*
  - *if $w^\bullet$ admits an initial cycle state $u^\bullet$ such that $\mathtt{pr}(u^\bullet) = v$, then the pair $(w^\bullet, u^\bullet)$ belongs to $R^\bullet$ and is associated with a back edge.*

As mentioned earlier, knowing which edges are forward edges or back edges, uniquely determines a tree decomposition. We denote by $(\mathscr{T}_0(\mathscr{K}), \mathsf{f}(\mathscr{K}))$ the tree decomposition associated with the above definition.

Note that ε is the only state of $\mathscr{U}_{\mathscr{K}}$ that does not admit any initial cycle state. It follows from the definition of $R^\bullet$ that all the successors of ε in $\mathscr{U}_{\mathscr{K}}$ are of the form $(w, \mathtt{new})$ (where w is a successor in $\mathscr{K}$ of the initial state of $\mathscr{K}$). Intuitively, the tree with back edges $\mathscr{U}_{\mathscr{K}}$ is defined as follows. We consider the usual unwinding construction [4] of a Kripke model and we modify it in two steps. First, in the unwinding construction, given a track ρ with $\mathsf{lst}(\rho) = w$ and given a pair (w, v) in the transition relation R, we construct *one successor* of ρ of the form $\rho \cdot v$. Here, we make two "copies" of the successor $\rho \cdot v$, one labeled by $\mathtt{new}$ and the other one labeled by $\mathtt{cycle}$.

[4] That is, the Kripke model with domain $\{\rho : \rho \text{ is a track in } \mathscr{K}\}$, initial state ε, transition relation $\{(\rho, \rho w) : \rho \text{ and } \rho \cdot w \text{ are tracks in } \mathscr{K}\}$ and a labeling function mapping each track ρ to the set $\mathsf{Llst}(\rho)$.

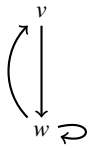

Figure 1: The Kripke Model $\mathscr{K}_0$.

The second modification is as follows: we delete certain edges and replace them with back edges (and finally, delete all the states that are reachable from ε). An edge from track ρ_1 to ρ_2 is deleted iff ρ_2 is labeled by `cycle` and ρ_2 and the initial cycle state of ρ_1 are labeled by the same state of $\mathscr{K}$.

In order to illustrate the construction $\mathscr{U}_\mathscr{K}$, we provide an example in Figure 1 and Figure 2. To make notation easier in the figure, we abbreviate `new` by `n` and `cycle` by `c`. Also, instead of writing ρ for a state, we only write the pair of labels $\mathsf{lst}(\rho)$. The back edges are those that are not straight lines.

Theorem 3. $\mathrm{CTL}^\star_\circlearrowleft$ *has a tree-like model property. Every satisfiable formula of* $\mathrm{CTL}^\star_\circlearrowleft$ *is satisfiable in a tree with back edges.*

This follows immediately from the following proposition.

Proposition 1. *Let* $\mathscr{K} = \langle \mathrm{AP}, \mathrm{W}, R, \mathsf{L}, w_I \rangle$ *be a Kripke model, let* $\mathscr{U}_\mathscr{K} = \langle \mathrm{AP}^\bullet, \mathrm{W}^\bullet, R^\bullet, \mathsf{L}^\bullet, w\bullet_I \rangle$ *be its tree-like unwinding. Then the relation* $\{(w^\bullet, \mathtt{pr}(w^\bullet)) : w^\bullet \in \mathrm{W}^\bullet\}$ *is a cycle-bisimulation. Hence,* $\mathscr{K}$ *and* $\mathscr{U}_\mathscr{K}$ *satisfy exactly the same formulas in* $\mathrm{CTL}^\star_\circlearrowleft$.

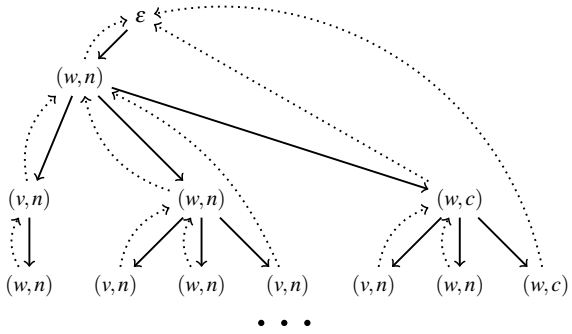

Figure 2: The tree with back edges $\mathscr{T}(\mathscr{K}_0)$.

Before finishing the section on model properties, we state one more property concerning the tree-like unwinding of a model. It states that if a formula is true in a tree-like unwinding, then we may assume the "witness" cycles (for the subformulas of the form $\mathrm{E}^\circlearrowleft_s \psi$) to be simple cycles (defined below). The property is not that interesting in itself, but it will play an important role in the next section for obtaining a 2ExpTime upper-bound for the satisfiability problem of the existential fragment of $\mathrm{CTL}^\star_\circlearrowleft$.

Definition 7. *A cycle* π *is a* simple cycle *if there is a sequence* $(n_i)_{i\in\mathbb{N}}$ *such that*

- $n_i < n_{i+1}$ *and* $\pi_{n_i} = \pi_0$, *for all* $i \in \mathbb{N}$;
- *for all* $i \in \mathbb{N}$ *and for all* $n_i < j < k < n_{i+1}$, *we have* $\pi_j \neq \pi_k$.

A (state or path) formula φ *is in* normal form *if for all subformulas* $\neg\psi$ *in* $\mathsf{sub}(\varphi)$, *the formula* ψ *is a variable. Given a formula* φ *in normal form, we define its* simple cycle translation *as the formula obtained by replacing each symbol* $\mathrm{E}^\circlearrowleft$ *in the formula* φ, *by the symbol* $\mathrm{E}^\circlearrowleft_s$. *The simple cycle translation of* φ *is denoted by* $(\varphi)_s$.

The semantics of the formulas of the form $(\varphi)_s$ *is defined by induction on* φ. *The basic and induction cases are defined as in Definition 3, with the additional induction step:* $\mathscr{K}, w \models \mathrm{E}^\circlearrowleft_s \psi$ *if there is a* **simple** *cycle* π *with beginning state w, such that* $\mathscr{K}, \pi \models \psi$.

Proposition 2. *Let* φ *be a formula in* $\mathrm{CTL}^\star_\circlearrowleft$ *in normal form and let* $\mathscr{K}$ *be a Kripke model. Then* $\mathscr{K} \models \varphi$ *iff* $\mathscr{U}_\mathscr{K} \models (\varphi)_s$, *where* $\mathscr{U}_\mathscr{K}$ *is the tree with back edges as in Definition 5.*

Expressiveness We now investigate the expressive power of $\mathrm{CTL}^\star_\circlearrowleft$ *w.r.t.* the usual temporal logics. All the results are collected in the following theorem.

Theorem 4 (Expressiveness comparison)**.** $\mathrm{CTL}^\star_\circlearrowleft$ *is strictly more expressive than* $\mathrm{CTL}^\star$ *and is incomparable with the* μCalculus.

Proof. We observed in the proof of Theorem 1 that $\varphi_1 = \mathtt{AG}\neg\mathtt{E}^{\circlearrowleft}\top$ is satisfiable but does not admit any finite model. Since CTL* and the μCALCULUS have the finite-model property, this implies that φ_1 is not equivalent to any formula in CTL* or in the μCALCULUS.

By using the result that there is no LTL formula expressing that a proposition p is true in every even state [28], we can show that the μCALCULUS formula $\psi = \nu x.p \wedge \Box\Box x$ is not equivalent to any formula in CTL$^\star_{\circlearrowleft}$. Note that ψ is true in a model if for all paths π starting from the initial state, p is true in every even state $(\pi)_{2i}$ of the path π. □

4 Decision Problems

In this section, we deal with the solution of the model-checking and satisfiability problems for CTL$^\star_{\circlearrowleft}$. Regarding the former, we show that we retain the same complexity as for CTL*, that is PSPACE. Concerning satisfiability, we also retain the same complexity of CTL* if we restrict to the existential-cycle fragment of the logic, that is 2EXPTIME. Conversely, we show that it is 3EXPTIME for the whole logic.

Model Checking For the solution of the model-checking problem of CTL$^\star_{\circlearrowleft}$, we employ a standard bottom-up procedure on the nesting of the path quantifiers of the specification under exam, which extends the one originally proposed for CTL* [10]. With more details, starting from the innermost state formulas φ of the kind $\mathtt{E}\psi$, $\mathtt{A}\psi$, $\mathtt{E}^{\circlearrowleft}\psi$, and $\mathtt{A}^{\circlearrowleft}\psi$, we determine their truth value over a KS $\mathscr{K}$ at a world $w \in \mathrm{W}$ by checking the emptiness of a suitable nondeterministic Büchi word automaton $\mathscr{N}^{\varphi}_{\mathscr{K},w}$. In case of a positive result, we enrich the labeling of the world w with a fresh proposition φ representing the formula φ itself. Obviously, the path formula ψ is just seen as a classic LTL formula, where all its subformulas of the kind described above are interpreted as atomic propositions whose truth values on the worlds of $\mathscr{K}$ are already computed in some previous step of the algorithm. It is important to observe that the difference between the automata for $\mathtt{E}\psi$ or $\mathtt{A}\psi$ and those for $\mathtt{E}^{\circlearrowleft}\psi$ or $\mathtt{A}^{\circlearrowleft}\psi$ resides in the fact that, for the latter, we have to further verify that the initial state of the path is seen infinitely often. This can be done by means of the standard Büchi acceptance condition. Hence, we directly obtain that the model checking for CTL$^\star_{\circlearrowleft}$ is not more complex than the same problem for CTL*.

Theorem 5. *The model-checking problem for* CTL$^\star_{\circlearrowleft}$ *is* PSPACE-COMPLETE w.r.t. *the formula complexity and* NLOGSPACE-COMPLETE w.r.t. *the data complexity.*

Satisfiability Differently from the model checking, the two introduced looping quantifiers $\mathtt{E}^{\circlearrowleft}\psi$ and $\mathtt{A}^{\circlearrowleft}\psi$ heavily affect the satisfiability of CTL$^\star_{\circlearrowleft}$. In particular, since this logic lacks of the standard tree-model property, we cannot use, for the CTL* part of CTL$^\star_{\circlearrowleft}$, the automata approach as proposed in [21]. Instead, we use symmetric two-way alternating tree automata [5], a simplified version of two-way graded alternating parity tree automata [5], which simply extend classic two-way alternating automata over ranked trees [26] to "unranked trees", *i.e.*, trees with possibly unbounded width. These are automata that allow to traverse a tree in forward and backward. We use these automata here to search for tree representation of the tree-like unwinding of a structure, as described in the previous section. With more details, for every CTL$^\star_{\circlearrowleft}$ state formula φ, we build an alternating parity two-way tree automaton $\mathscr{A}_\varphi$ such that a KS $\mathscr{K} = \langle \mathrm{AP}, \mathrm{W}, R, \mathsf{L}, w_I\rangle$ is a model of φ iff $\mathscr{A}_\varphi$ accepts a tree $\mathscr{T}_\mathscr{K} = \langle \mathrm{AP} \cup \{\mathtt{new}, \uparrow\}, \mathrm{W}^\bullet, R^\star, \mathsf{L}^\star, w_I^\bullet\rangle$ associated with the tree-like unwinding $\mathscr{U}_\mathscr{K} = \langle \mathrm{AP}, \mathrm{W}^\bullet, R^\bullet, \mathsf{L}^\bullet, w_I^\bullet\rangle$ of $\mathscr{K}$ via the following properties: *(i)* $R^\star = \{(w^\bullet, v^\bullet) \in R^\bullet : |w^\bullet| < |v^\bullet|\}$, *(ii)* $\mathsf{L}^\star(w^\bullet) \cap \mathrm{AP} = \mathsf{L}^\bullet(w^\bullet)$, *(iii)* $\mathtt{new} \in \mathsf{L}^\star(w^\bullet)$ iff $\mathsf{lst}(w^\bullet) = (w, \mathtt{new})$, for some $w \in \mathrm{W}$, and *(iv)* $\uparrow \in \mathsf{L}^\star(w^\bullet)$ iff there exists $(w^\bullet, v^\bullet) \in R^\bullet$ with $|v^\bullet| < |w^\bullet|$. Intuitively, $\mathscr{T}_\mathscr{K}$ is built from $\mathscr{U}_\mathscr{K}$ by deleting all back edges (property *(i)*) and enriching the original labeling of every

world $w^\bullet$ (property *(ii)*) with `new`, if the last letter of $w^\bullet$ contains the flag with the same name (property *(iii)*), and with $\uparrow$, if $w^\bullet$ is the origin of a back edge (property *(iv)*). It is not hard to see that, for every unwinding $\mathscr{U}_\mathscr{K}$ of a KS $\mathscr{K}$, there exists one and only one tree $\mathscr{T}_\mathscr{K}$ satisfying the previous four properties. Therefore, instead of looking for a model $\mathscr{K}$ of φ or its tree-like unwinding $\mathscr{U}_\mathscr{K}$, we just look for its tree representation $\mathscr{T}_\mathscr{K}$. This idea is at the basis for the automata-theoretic approach described in the proofs of the following theorems.

Theorem 6. *The satisfiability problem for* $\mathrm{CTL}^\star_{\circlearrowleft}$ *can be solved in 3*ExpTime *and is 2*ExpTime-hard.

Proof. The 2ExpTime lower bound for $\mathrm{CTL}^\star_{\circlearrowleft}$ immediately follows from the one of $\mathrm{CTL}^\star$. For the 3ExpTime upper bound, given a $\mathrm{CTL}^\star_{\circlearrowleft}$ state formula φ, we reduce the associated satisfiability question to the emptiness problem of an alternating parity two-way tree automaton $\mathscr{A}_\varphi$, whose size and index are, respectively, doubly and single exponential in $|\varphi|$. For a detailed definition of symmetric alternating parity two-way tree automata and the related concepts of size and index, we refer to [5]. Since the emptiness of $\mathscr{A}_\varphi$ can be checked in time exponential *w.r.t.* both its states and index [5], we obtain the desired result [5].

As mentioned above, $\mathscr{A}_\varphi$ needs to recognize all and only the tree representations $\mathscr{T}_\mathscr{K}$ of the tree-like unwindings $\mathscr{U}_\mathscr{K}$ of KS models $\mathscr{K}$ of φ. As it is usually done for $\mathrm{CTL}^\star$, we slightly weaken this property by allowing $\mathscr{A}_\varphi$ to run on trees that also contain, as labeling of its worlds, the subformulas of φ of the form $\mathtt{E}\psi$, $\mathtt{A}\psi$, $\mathtt{E}^{\circlearrowleft}\psi$, and $\mathtt{A}^{\circlearrowleft}\psi$, which are interpreted as fresh atomic propositions. We denote by $\mathsf{sub}_Q(\varphi)$ the set of subformulas of φ of the form $\mathtt{E}\psi$, $\mathtt{A}\psi$, $\mathtt{E}^{\circlearrowleft}\psi$, and $\mathtt{A}^{\circlearrowleft}\psi$. We also let $\mathsf{sub}^\neg_Q(\varphi)$ be the the closure under negation of the set $\mathsf{sub}_Q(\varphi)$, *i.e.*, for every $\mathtt{E}\psi$ (resp., $\mathtt{A}\psi$, $\mathtt{E}^{\circlearrowleft}\psi$, $\mathtt{A}^{\circlearrowleft}\psi$) in $\mathsf{sub}(\varphi)$, we have $\mathtt{A}\neg\psi$ (resp., $\mathtt{E}\neg\psi$, $\mathtt{A}^{\circlearrowleft}\neg\psi$, $\mathtt{E}^{\circlearrowleft}\neg\psi$) in $\mathsf{sub}^\neg(\varphi)$. So, instead of considering a model $\mathscr{K} = \langle \mathrm{AP}, \mathrm{W}, R, \mathsf{L}, w_I\rangle$ of φ, we work on the enriched KS $\mathscr{K}^\star = \langle \mathrm{AP}^\star, \mathrm{W}, R, \mathsf{L}^\star, w_I\rangle$ such that *(i)* $\mathrm{AP}^\star = \mathrm{AP} \cup \mathsf{sub}^\neg(\varphi)$, *(ii)* $\mathsf{L}^\star(w) \cap \mathrm{AP} = \mathsf{L}(w)$, and *(iii)* $\eta \in \mathsf{L}^\star(w)$ iff $\mathscr{K}, w \models \eta$, for all $\eta \in \mathsf{sub}^\neg(\varphi)$. set $\mathsf{sub}(\varphi)$, *i.e.*, for every $\mathtt{E}\psi$ (resp., $\mathtt{A}\psi$, $\mathtt{E}^{\circlearrowleft}\psi$, $\mathtt{A}^{\circlearrowleft}\psi$) in $\mathsf{sub}(\varphi)$, we have $\mathtt{A}\neg\psi$ (resp., $\mathtt{E}\neg\psi$, $\mathtt{A}^{\circlearrowleft}\neg\psi$, $\mathtt{E}^{\circlearrowleft}\neg\psi$) in $\mathsf{sub}^\neg(\varphi)$.

The automaton $\mathscr{A}_\varphi$ is built as the conjunction of an automaton $\mathscr{A}_\eta$, for every subformula $\eta \in \mathsf{sub}^\neg(\varphi)$, and a deterministic safety (*i.e.*, without acceptance condition) automaton $\mathscr{D}_\varphi$ used to verify that φ is satisfied at the root of the input tree $\mathscr{T}_\mathscr{K}$, when φ is interpreted as a Boolean formula on $\mathrm{AP}^\star$. In addition, $\mathscr{A}_\varphi$ needs to check that, if a world is not labeled by a state formula $\eta \in \mathsf{sub}^\neg(\varphi)$, it is necessarily labeled by a formula $\overline{\eta} \in \mathsf{sub}^\neg(\varphi)$ equivalent to its negation, *i.e.*, $\overline{\eta} \equiv \neg\eta$. The automaton $\mathscr{A}_\eta$ is committed to check that a world labeled by $\eta \in \mathsf{sub}^\neg(\varphi)$ really satisfies this formula. Formally, we have $\mathscr{A}_\varphi \triangleq \mathscr{D}_\varphi \wedge \bigwedge_{\eta\in\mathsf{sub}^\neg(\varphi)} \mathscr{A}_\eta$. So, its size is the sum of the sizes of the components. The construction of $\mathscr{D}_\varphi$ is trivial. Moreover, the automata for $\mathtt{A}\psi$ and $\mathtt{A}^{\circlearrowleft}\psi$ can be directly derived from the automaton for $\mathtt{E}\neg\psi$ and $\mathtt{E}^{\circlearrowleft}\neg\psi$ by replacing $\vee$ and $\Diamond$ with $\wedge$ and $\Box$ in their definitions. Hence, we just focus on the constructions for the latter.

We start with the construction of $\mathscr{A}_{\mathtt{E}\psi}$ for $\mathtt{E}\psi$. Consider the nondeterministic Büchi word automaton $\mathscr{N}_\psi = \langle 2^{\mathrm{AP}^\star}, \mathrm{Q}, \delta, \mathrm{Q}_I, \mathrm{F}\rangle$ obtained by applying the Vardi-Wolper construction to ψ which is read as an LTL formula over $\mathrm{AP}^\star$ [27]. We set as a two-way Büchi tree automaton $\mathscr{A}_{\mathtt{E}\psi} \triangleq \langle \Sigma^\star, \mathrm{Q}^\star, \delta^\star, q_I^\star, \mathrm{F}^\star\rangle$, where the alphabet $\Sigma^\star \triangleq 2^{\mathrm{AP}^\star \cup \{\mathtt{new}, \uparrow\}}$ augments the set of extended atomic propositions $\mathrm{AP}^\star$ with the symbols `new` and $\uparrow$, as required by the definition of the tree representations $\mathscr{T}_\mathscr{K}$. The set of states $\mathrm{Q}^\star \triangleq \{q_I^\star\} \cup \mathrm{Q} \times \{\downarrow, \uparrow\}$ contains the initial state $q_I^\star$ plus two copies of the states of $\mathscr{N}_\psi$, one for each direction of navigation over the tree $\mathscr{T}_\mathscr{K}$. For the Büchi acceptance condition we consider the set $\mathrm{F}^\star \triangleq \{q_I^\star\} \cup \mathrm{F} \times \{\downarrow\}$. The definition of the transition function $\delta^\star$ follows. For the sake of readability, we divide it in three parts, depending on whether it predicates on $q_I^\star$, a state q flagged with $\downarrow$ or a state q flagged with $\uparrow$.

[5] In particular, Theorem 6.7 in [5] can be used for the translation. Observe that, since we do not make use of any graded modalities (our box and diamond symbols stand for $[\![0]\!]$ and $\langle\!\langle 0\rangle\!\rangle$ in their syntax) the resulting automaton is simply a symmetric non-deterministic tree automaton.

- The initial state $q_I^\star$ is used to start the evaluation of the formula $\mathtt{E}\psi$ on every world of the input tree labeled by $q_I^\star$. This is done by starting the simulation of $\mathscr{N}_\psi$. Formally, we have that $\delta^\star(q_I^\star,\sigma) \triangleq (\Box, q_I^\star) \wedge \bigvee_{q\in \mathrm{Q}_I}(\varepsilon,(q,\downarrow))$, if $\mathtt{E}\psi \in \sigma$, and $\delta^\star(q_I^\star,\sigma) \triangleq (\Box, q_I^\star)$, otherwise.
- Every copy of a state $q \in \mathrm{Q}$ flagged with $\downarrow$ is used to effectively verify the existence of an infinite path in $\mathscr{U}_\mathscr{K}$ satisfying ψ. This is done by guessing an extension of the finite path built up to now and sending, to the corresponding direction, a successor p of q that complies with the transition function δ of $\mathscr{N}_\psi$, when the labeling σ of the world under exam is read. As the input tree $\mathscr{T}_\mathscr{K}$ is a representation of the tree with back edges $\mathscr{U}_\mathscr{K}$, we have also to take them into account when we guess the extension of the path from a world labeled with $\uparrow$. This is done by sending up along the tree the copy of the state p flagged with $\uparrow$, which is used to simulate a jump to the world destination of the back edge. Formally, we have $\delta^\star((q,\downarrow),\sigma) \triangleq \bigvee_{p\in\delta(q,\sigma\cap \mathrm{AP}^\star)}(\Diamond,(p,\downarrow)) \vee \uparrow(p)$, where $\uparrow(p)$ is set to $(\varepsilon,(p,\uparrow))$ if $\uparrow \in \sigma$, and to $\mathtt{f}$, otherwise.
- Finally, for every copy of a state $q \in \mathrm{Q}$ flagged with $\uparrow$, we only have to modify the state and the direction of the automaton when we are approaching to the destination of the back edge that gave rise to the evaluation of $(q,\uparrow)$. Fortunately, due to the structure of the tree-like unwinding $\mathscr{U}_\mathscr{K}$ and, consequently, of its tree representation $\mathscr{T}_\mathscr{K}$, when we reach a world labeled by $\mathtt{new}$, we are sure that the immediate ancestor of this world is the destination of the back edge. Thus, we can immediately change the flag of the state q to $\downarrow$ in order to resume the verification of the path formula ψ. Formally, $\delta^\star((q,\uparrow),\sigma) \triangleq (\uparrow,(q,\downarrow))$, if $\mathtt{new} \in \sigma$, and $\delta^\star((q,\uparrow),\sigma) \triangleq (\uparrow,(q,\uparrow))$, otherwise.

Now, by construction, it is not hard to prove that $\mathscr{A}_{\mathtt{E}\psi}$ correctly verifies that every world of $\mathscr{T}_\mathscr{K}$ labeled by $\mathtt{E}\psi$ satisfies $\mathtt{E}\psi$ in $\mathscr{U}_\mathscr{K}$. Also, by the Vardi-Wolper procedure, it follows that $|\mathrm{Q}| = \mathrm{O}(2^{|\psi|})$. Consequently, the size of $\mathscr{A}_{\mathtt{E}\psi}$ is exponential in the length of $\mathtt{E}\psi$.

The construction of $\mathscr{A}_{\mathtt{E}^\circlearrowleft\psi}$ is quite more complex than the one previously described, as it also requires a projection operation that is the reason behind the exponential gap between the upper and lower bounds. Differently from the automata for classic path quantifiers, we cannot evaluate the correctness of the labeling $\mathtt{E}^\circlearrowleft\psi$ on all worlds of the tree in one shot. This is because of the possible interactions among the cycles starting in different worlds, which does not allow us to determine which is the origin of the path we are interested in. Consequently, we have to focus on one world labeled by $\mathtt{E}^\circlearrowleft\psi$ at a time and check the existence of a path passing infinitely often through that world, which also satisfies the property ψ. This unique world is identified by a fresh symbol #. Then, an universal projection operation over such a symbol will take care of the fact that this check has to be done for every possible world labeled by $\mathtt{E}^\circlearrowleft\psi$. Formally, $\mathscr{A}_{\mathtt{E}^\circlearrowleft\psi}$ is built as follows: $\Pi^\forall_\#(\mathscr{N}_\# \vee \mathscr{A}^\#_{\mathtt{E}^\circlearrowleft\psi})$. Intuitively, we make a universal projection over # of a disjunction between the automaton $\mathscr{N}_\#$, accepting all trees where the labeling # is incorrect (*i.e.*, there are more than one occurrences of # or this symbol is on a world that is not labeled by $\mathtt{E}^\circlearrowleft\psi$), and the automaton $\mathscr{A}^\#_{\mathtt{E}^\circlearrowleft\psi}$, verifying the existence of a path satisfying ψ that starts and passes infinitely often through the world labeled by #. The construction of $\mathscr{N}_\#$ is trivial. For the computation of the projection, we use the equality $\Pi^\forall_\#\mathscr{A} = \neg\Pi^\exists_\#\neg\mathscr{A}$. Note however that there is no known projection operation that can act directly on a two-way automaton. Instead, we have first to translate it into a nondeterministic one-way automaton [5] and then apply the standard projection. Due to the nondeterminization procedure, $\Pi^\forall_\#\mathscr{A}$ has exponential size *w.r.t.* that of $\mathscr{A}$. So, $\mathscr{A}_{\mathtt{E}^\circlearrowleft\psi}$ is exponential in the size of $\mathscr{A}^\#_{\mathtt{E}^\circlearrowleft\psi}$.

It remains to define the latter automaton. As above, let $\mathscr{N}_\psi = \langle 2^{\mathrm{AP}^\star}, \mathrm{Q}, \delta, \mathrm{Q}_I, \mathrm{F}\rangle$ be the nondeterministic Büchi word automaton obtained by applying the Vardi-Wolper construction to ψ. Then, we set $\mathscr{A}^\#_{\mathtt{E}^\circlearrowleft\psi} \triangleq \langle \Sigma^\star, \mathrm{Q}^\star, \delta^\star, q_I^\star, \mathrm{F}^\star\rangle$ as a two-way Büchi tree automaton having alphabet $\Sigma^\star \triangleq 2^{\mathrm{AP}^\star \cup \{\mathtt{new},\uparrow,\#\}}$. The set of states $\mathrm{Q}^\star \triangleq \{q_I^\star\} \cup \mathrm{Q} \times \{\mathtt{f},\mathtt{t}\} \times \{\#,\downarrow,\uparrow\}$ contains the initial state $q_I^\star$ plus six copies of the states of $\mathscr{N}_\psi$.

Each of them is flagged with a Boolean value keeping track of the original acceptance condition derived from $\mathscr{N}_\psi$ and a symbol indicating the direction of navigation over the tree. Differently from the previous case, we have also # as a flag in order to indicate the passage though the state labeled by the flag itself. For the Büchi acceptance condition we consider the set $\mathrm{F}^\star \triangleq \{q_I^\star\} \cup \mathrm{Q} \times \{\mathtt{t}\} \times \{\#\}$. Intuitively, apart from the initial state, we assume as final those states that certify both the passage through the origin of the path indicated by # and the possibly previous occurrence of an accepting state. It remains to define the transition function $\delta^\star$. Here we use $\alpha(q,\alpha)$ to denote the Boolean value $\mathtt{t}$, if $q \in \mathrm{F}$, and α, otherwise.

- The initial state $q_I^\star$ is used to start evaluating the formula $\mathtt{E}^{\circlearrowleft}\psi$ on the unique world of the input tree labeled by #. Formally, we have $\delta^\star(q_I^\star,\sigma) \triangleq \bigvee_{q\in \mathrm{Q}_I}(\varepsilon,(q,\alpha(q,\mathtt{f}),\downarrow))$, if $\# \in \sigma$, and $\delta^\star(q_I^\star,\sigma) \triangleq (\Box, q_I^\star)$, otherwise. Note that, since we are just starting with the simulation of $\mathscr{N}_\psi$, the flag $\alpha(q,\mathtt{f})$ concerning the memory on the acceptance condition only depends on the state q, as the second argument is fixed to $\mathtt{f}$.
- Since a state $(q,\alpha,\#)$ is simply used to verify the passage through the starting point of the path satisfying ψ, the automaton has to reset the memory on the acceptance condition and continue with the simulation of $\mathscr{N}_\psi$. Formally, $\delta^\star((q,\alpha,\#),\sigma) \triangleq (\varepsilon,(q,\mathtt{f},\downarrow))$.
- The automaton $\mathscr{A}^{\#}_{\mathtt{E}^{\circlearrowleft}\psi}$ on the state $(q,\alpha,\downarrow)$ behaves similar to $\mathscr{A}_{\mathtt{E}\psi}$ on $(q,\downarrow)$. One difference resides in the update $\alpha(p,\alpha)$ of the memory on the acceptance condition, which takes into account both the previous memory α and the membership of p in F. The other difference is that, if σ contains the symbol #, we have to record this fact in the state, by swapping the flag from $\downarrow$ to #. Formally, we have $\delta^\star((q,\alpha,\downarrow),\sigma) \triangleq \bigvee_{p\in\delta(q,\sigma\cap \mathrm{AP}^\star)}(\Diamond,(p,\alpha(p,\alpha),\beta)) \vee \uparrow(p)$, where $\beta = \#$, if $\# \in \sigma$, and $\beta = \downarrow$, otherwise; moreover, $\uparrow(p)$ is set to $(\varepsilon,(p,\alpha(p,\alpha),\uparrow))$ if $\uparrow \in \sigma$, and to $\mathtt{f}$, otherwise.
- Finally, as for $\mathscr{A}_{\mathtt{E}\psi}$, a state of the form $(q,\alpha,\uparrow)$ identifies the destination of a back edge. Thus, we have $\delta^\star((q,\alpha,\uparrow),\sigma) \triangleq (\uparrow,(q,\alpha,\downarrow))$, if $\mathtt{new} \in \sigma$, and $\delta^\star((q,\alpha,\uparrow),\sigma) \triangleq (\uparrow,(q,\alpha,\uparrow))$, otherwise.

Finally, the size of $\mathscr{A}^{\#}_{\mathtt{E}^{\circlearrowleft}\psi}$ is exponential in the length of $\mathtt{E}^{\circlearrowleft}\psi$, which implies that $\mathscr{A}_{\mathtt{E}^{\circlearrowleft}\psi}$ is doubly exponential *w.r.t.* the same length. □

In case we want to restrict our attention to the satisfiability of the $\mathrm{CTL}^\star_{\circlearrowleft}$ fragment having only existential looping quantifiers, we can improve the previous proof, obtaining a tight 2ExpTime procedure, by providing a single exponential construction for the automaton $\mathscr{A}_{\mathtt{E}^{\circlearrowleft}\psi}$. Indeed, thanks to the simple cycle property of the verification of the formula $\mathtt{E}^{\circlearrowleft}\psi$ on the tree-like unwinding $\mathscr{U}_\mathscr{K}$, we can just focus on cycle paths of $\mathscr{U}_\mathscr{K}$ going through the successors of their origin labeled by $\mathtt{new}$. In this way, there are no interactions among the paths that start at different worlds labeled by $\mathtt{E}^{\circlearrowleft}\psi$, since two paths passing through the same world necessarily use different successors. Consequently, we can always uniquely identify the origin of a path on which we have to pass infinitely often.

Unfortunately, the same idea cannot be exploited for the verification of the universal looping quantifiers $\mathtt{A}^{\circlearrowleft}\psi$, as we have to check the property ψ on all cycle paths and not only on those that are simple. At the moment, it is left open whether a 2ExpTime satisfiability procedure for the whole $\mathrm{CTL}^\star_{\circlearrowleft}$ logic exists.

Theorem 7. *The satisfiability problem for the existential-cycle fragment of* $\mathrm{CTL}^\star_{\circlearrowleft}$ *is* 2ExpTime-complete.

5 Discussion

To conclude, we give a concise overview of the main properties of the cycle-logic extension we have introduced. We also explain why this extension is natural and why, given our results, we have decided to focus our presentation on the cycle-logic $\mathrm{CTL}^\star_{\circlearrowleft}$.

$\text{CTL}^\star_{\circlearrowleft}$ allows us to quantify only over cycles, that is, paths such that the initial state occurs infinitely often. Hence, it has been natural to consider also a more general logic (denoted by $\text{ECTL}^\star_{\circlearrowleft}$) allowing us to test whether *any arbitrary* state of a given path occurs infinitely often in the path. More formally, $\text{ECTL}^\star_{\circlearrowleft}$ is the extension of $\text{CTL}^\star$ obtained by adding the symbol $\circlearrowleft$. This symbol is treated as an atomic path formula and is true at a state in a path iff the state occurs infinitely often in the path. It is easy to see that $\text{ECTL}^\star_{\circlearrowleft}$ is an extension of $\text{CTL}^\star_{\circlearrowleft}$.

We have studied several properties about $\text{ECTL}^\star_{\circlearrowleft}$ and, among the others, we have shown that this logic does not preserve the cycle-bisimulation property. Clearly, one can use a stronger notion of bisimulation under which $\text{ECTL}^\star_{\circlearrowleft}$ can still retain the invariance. However, we came up with notions that are not very intuitive (as the notion of cycle-bisimulation) and not useful to prove any kind of tree-like model property. Given these negative results, we decided to not present extensively this part.

We would like to mention that we also considered the extension of LTL with the symbol $\circlearrowleft$. It is not hard to show that it is a proper extension of LTL and is orthogonal to ω-regular expressions. We can also prove that the finite satisfiability problem for that logic is decidable (using an adaption of the proof for LTL [27]). We did not present these results by lack of space.

Finally, as future work we would like to investigate the use of the introduced cycle construct in the realm of logics for multi-agent systems such as $\text{ATL}^\star$ [2] and Strategy Logic [23]. These logics have been proved to be useful to reasoning about strategic abilities in a number of complicated settings. In particular, the latter is able to express sophisticated solution concepts such as Nash Equilibria and Subgame Perfect Equilibria, as well as they it has been used to express iterative extensive game forms such as the iterated prisoner dilemma. In all these contexts, talking explicitly about cycles could play a central role in solving the related game questions.

Acknoledgments

Aniello Murano and Loredana Sorrentino are partially supported by the GNCS 2016 project: Logica, Automi e Giochi per Sistemi Auto-adattivi. Giuseppe Perelli thanks the support of the ERC Advanced Grant 291528 (“Race”) at Oxford.

References

[1] R. Alur & T.A. Henzinger (1998): *Finitary Fairness.* *TOPLAS* 20(6), pp. 1171–1194, doi:10.1145/295656.295659.

[2] R. Alur, T.A. Henzinger & O. Kupferman (2002): *Alternating-Time Temporal Logic.* *JACM* 49(5), pp. 672–713, doi:10.1145/585265.585270.

[3] R. Alur & S. La Torre (2004): *Deterministic generators and games for LTL fragments.* *ACM Transactions on Computational Logic (TOCL)* 5(1), pp. 1–25, doi:10.1145/963927.963928.

[4] B. Aminof, A. Murano, S. Rubin & F. Zuleger (2016): *Prompt Alternating-Time Epistemic Logics.* In: *KR'16*, AAAI Press, pp. 258–267.

[5] P.A. Bonatti, C. Lutz, A. Murano & M.Y. Vardi (2008): *The Complexity of Enriched muCalculi.* *LMCS* 4(3), pp. 1–27, doi:10.2168/LMCS-4(3:11)2008.

[6] L. Bozzelli, A. Murano & A. Peron (2010): *Pushdown Module Checking.* *FMSD* 36(1), pp. 65–95, doi:10.1007/s10703-010-0093-x.

[7] K. Chatterjee, T.A. Henzinger & F. Horn (2010): *Finitary Winning in omega-Regular Games.* *TOCL* 11(1), pp. 1:1–26, doi:10.1145/1614431.1614432.

[8] E.M. Clarke & E.A. Emerson (1981): *Design and Synthesis of Synchronization Skeletons Using Branching-Time Temporal Logic.* In: *LP'81*, LNCS 131, Springer, pp. 52–71, doi:10.1007/BFb0025774.

[9] E.M. Clarke, E.A. Emerson & A.P. Sistla (1986): *Automatic Verification of Finite-State Concurrent Systems Using Temporal Logic Specifications.* *TOPLAS* 8(2), pp. 244–263, doi:10.1145/5397.5399.

[10] E.M. Clarke, O. Grumberg & D.A. Peled (2002): *Model Checking.* MIT Press.

[11] E.A. Emerson & J.Y. Halpern (1985): *Decision Procedures and Expressiveness in the Temporal Logic of Branching Time.* *JCSS* 30(1), pp. 1–24, doi:10.1016/0022-0000(85)90001-7.

[12] E.A. Emerson & J.Y. Halpern (1986): *"Sometimes" and "Not Never" Revisited: On Branching Versus Linear Time.* *JACM* 33(1), pp. 151–178, doi:10.1145/4904.4999.

[13] E.A. Emerson & C.S. Jutla (1988): *The Complexity of Tree Automata and Logics of Programs (Extended Abstract).* In: *FOCS'88*, IEEE Computer Society, pp. 328–337, doi:10.1109/SFCS.1988.21949.

[14] A. Ferrante, A. Murano & M. Parente (2008): *Enriched Mu-Calculi Module Checking.* *LMCS* 4(3), pp. 1–21, doi:10.2168/LMCS-4(3:1)2008.

[15] E. Grädel, W. Thomas & T. Wilke (2002): *Automata, Logics, and Infinite Games: A Guide to Current Research.* LNCS 2500, Springer, doi:10.1007/3-540-36387-4.

[16] S.A. Kripke (1963): *Semantical Considerations on Modal Logic.* *APF* 16, pp. 83–94, doi:10.1002/malq.19630090502.

[17] O. Kupferman, G. Morgenstern & A. Murano (2006): *Typeness for omega-Regular Automata.* *IJFCS* 17(4), pp. 869–884, doi:10.1142/S0129054106004157.

[18] O. Kupferman, N. Piterman & M.Y. Vardi (2002): *Pushdown Specifications.* In: *LPAR'02*, LNCS 2514, Springer, pp. 262–277, doi:10.1007/3-540-36078-6_18.

[19] O. Kupferman, N. Piterman & M.Y. Vardi (2009): *From Liveness to Promptness.* *FMSD* 34(2), pp. 83–103, doi:10.1007/s10703-009-0067-z.

[20] O. Kupferman, A. Pnueli & M.Y. Vardi (2012): *Once and For All.* *JCSS* 78(3), pp. 981–996, doi:10.1016/j.jcss.2011.08.006.

[21] O. Kupferman, M.Y. Vardi & P. Wolper (2000): *An Automata Theoretic Approach to Branching-Time Model Checking.* *JACM* 47(2), pp. 312–360, doi:10.1145/333979.333987.

[22] O. Kupferman, M.Y. Vardi & P. Wolper (2001): *Module Checking.* *IC* 164(2), pp. 322–344, doi:10.1006/inco.2000.2893.

[23] F. Mogavero, A. Murano, G. Perelli & M.Y. Vardi (2014): *Reasoning About Strategies: On the Model-Checking Problem.* *TOCL* 15(4), pp. 34:1–42, doi:10.1145/2631917.

[24] F. Mogavero, A. Murano & L. Sorrentino (2015): *On Promptness in Parity Games*. *Fundamenta Informaticae* 139(3), pp. 277–305, doi:10.3233/FI-2015-1235.

[25] A. Pnueli (1977): *The Temporal Logic of Programs.* In: *FOCS'77*, IEEE Computer Society, pp. 46–57, doi:10.1109/SFCS.1977.32.

[26] M.Y. Vardi (1998): *Reasoning about The Past with Two-Way Automata.* In: *ICALP'98*, LNCS 1443, Springer, pp. 628–641, doi:10.1007/BFb0055090.

[27] M.Y. Vardi & P. Wolper (1986): *An Automata-Theoretic Approach to Automatic Program Verification.* In: *LICS'86*, IEEE Computer Society, pp. 332–344.

[28] P. Wolper (1983): *Temporal Logic Can Be More Expressive.* *IC* 56(1-2), pp. 72–99, doi:10.1016/S0019-9958(83)80051-5.

[29] W. Zielonka (1998): *Infinite Games on Finitely Coloured Graphs with Applications to Automata on Infinite Trees.* *TCS* 200(1-2), pp. 135–183, doi:10.1016/S0304-3975(98)00009-7.